

# Dynamic Selection of Models for a Ventilator-Management Advisor

Geoffrey W. Rutledge, M.D.

Section on Medical Informatics, Department of Medicine

Stanford University, Stanford, California 94305-5479

rutledge@camis.stanford.edu

## ABSTRACT

*A ventilator-management advisor (VMA) is a computer program that monitors patients who are treated with a mechanical ventilator. A VMA implements a patient-specific physiologic model to interpret patient data and to predict the effects of alternative control settings for the ventilator. Because a VMA evaluates its physiologic model repeatedly during each cycle of data interpretation, highly complex models may require more computation time than is available in this time-critical application. On the other hand, less complex models may be inaccurate if they are unable to represent a patient's physiologic abnormalities. For each patient, a VMA should select a model that balances the tradeoff of prediction accuracy and computation-time complexity.*

*I present a method to select models that are at an appropriate level of detail for time-constrained decision tasks. The method is based on a local search in a graph of models (GoM) for a model that maximizes the tradeoff of computation-time complexity and prediction accuracy. For each model under consideration, a belief network computes a probability of model adequacy given the qualitative prior information, and the goodness of fit of the model to the data provides a measure of the conditional probability of adequacy given the quantitative observations.*

*I apply this method to the problem of model selection for a VMA. I describe an implementation of a graph of physiologic models that range in complexity from VentPlan, a simple model with 3 compartments, to VentSim, a multicompartment model with detailed airway, circulation and mechanical ventilator components. I demonstrate how the model-selection method selects dynamically, from this GoM, models that are likely to meet the accuracy requirement of a VMA.*

## THE VENTILATOR-MANAGEMENT ADVISOR MODELING PROBLEM

A ventilator-management advisor (VMA) is a computer program that monitors patients in an intensive-care unit (ICU) who are treated with a mechanical ventilator. A VMA implements a patient-specific physiologic model to interpret patient observations and to predict the effects of alternative proposed ventilator settings.

A VMA evaluates this model repeatedly during each cycle of fitting the model to the data and of searching for the computed optimal ventilator settings [15].

A VMA should incorporate prediction models that are accurate, yet tractable for use in this real-time patient-care application. A model that is highly complex requires longer computation time, and delays the data interpretation and treatment recommendation. On the other hand, a simplified model provides less accurate predictions of the effects of changes in the ventilator settings if it is unable to represent a specific patient's physiologic abnormalities. Ideally, a VMA would apply a model that balances the tradeoff of prediction accuracy and computation-time complexity for each patient it monitors.

## AUTOMATED MODEL SELECTION

There is a growing interest in methods to automate the tasks of creating and applying models that are at an appropriate level of detail [8]. One interesting approach is to organize a set of alternative models as a graph. In a *graph of models* (GoM), nodes represent models, and arcs represent the simplifying assumptions that distinguish adjacent models [11]. In the GoM formulation, a search algorithm explores the arcs from unsuitable models to adjacent models by evaluating the consequences of asserting or retracting the corresponding simplifying assumptions.

The assumptions for model selection in the GoM are categorical. In situations where all available models violate at least one assumption, there is no method to reason about the degree of assumption violation, or to find the model with the least severe violations. Prior investigators demonstrated that the GoM is a powerful tool to organize the search for a model that satisfies all modeling constraints. These investigators have not attempted to compare multiple models that satisfy all constraints, but rather stop the search when the first satisfying model is found [1,19,20].

## Optimal Model Selection under a Time Constraint

To apply the GoM formalism to the selection of models for a time-critical, model-based decision system, such as a VMA, requires that we compare a model's expected prediction accuracy with the computation-time delay that the model imposes. The optimal model

for a time-critical application represents a tradeoff of prediction accuracy and computation-time complexity.

I previously described a time-dependent, decision analytic definition of the optimal model to select under a time constraint. This definition is based on the integral over time of the value of alternative model selections [14]. Unfortunately, no definition of the optimal model is helpful for a real-time model-selection algorithm if we are unable to evaluate all models within the time available for model selection.

The GoM formalism provides a method to organize a set of alternative models, and to guide a heuristic search for a model that is *likely* to be optimal. The definition of the optimal model provides a means to assess (retrospectively) the performance of the heuristic search algorithm.

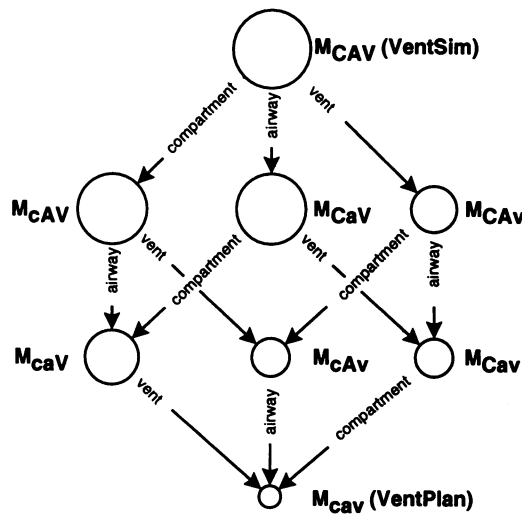


Figure 1. A graph of models for a ventilator-management advisor. Nodes represent models, and arcs represent the simplifying assumptions that separate the adjacent models. The models are arranged from the most complex model (VentSim) that makes no additional assumptions, to the least complex model (VentPlan) that makes three additional assumptions. More complex models are shown as larger circles. Upper-case and lower-case model subscripts indicate greater and lesser levels of detail in the corresponding components: C, c: circulation-compartment component; A, a: airway component; V, v: ventilator component.

#### Heuristic Model Selection under a Time Constraint

A local search within the GoM, for a model that is likely to be optimal, is guided by a search metric that combines a measure of a model's *cost* and a measure of the model's *benefit*.

**Cost of a Model Selection.** The cost of a model is a function of the computation-time delay that the model

imposes. This computation-time delay has two components:  $t_{model}$ , the time taken to fit the model to the observations (to make the model system-specific), and  $t_{search}$ , the time required to perform the search for the optimal control settings, once the system-specific model parameters are known.

An exponential function is useful to represent the cost of the time delay for time-critical control decisions [5]. I set

$$\text{Cost}(t) = C e^{t/k}, \quad (1)$$

in which  $t$  is the time delay relative to  $t = 0$ ,  $k$  is a time constant that specifies the time constraint, and  $C$  is a scaling constant. If a model is selected at  $t = 0$ , the cost of the model is  $\text{Cost}(t_{model} + t_{search})$ . Once a model is fitted to the data, the cost of generating that model's control recommendation is the incremental cost of  $t_{search}$ . The incremental cost of an additional time delay  $\delta t$ , after an initial time  $t_I$ , is  $\text{Cost}(t_I + \delta t) - \text{Cost}(t_I)$ , which increases as  $t_I$  increases.

**Benefit of a Model Selection.** The *benefit* of a model is a result of the value of its control recommendations. As the model-prediction accuracy increases, the control recommendations improve, and the value of the model increases. At some level of accuracy, further increases in accuracy do not change a model-based control-setting recommendation. That is, there is an accuracy, greater than which no increase in value occurs. This "maximum needed" prediction accuracy varies according to the domain, and depends on a complex interaction of the model predictions, control variable settings, and the loss function (or value model) used to optimize the control settings.

I define the event that a model  $M_i$  is *adequate* as

$$M_i^A = \bigcap_j [ \Pr(|\hat{y}_{ij} - y_{obsj}| < \epsilon \sigma_j) \geq 1 - \delta_A ], \quad (2)$$

in which  $\hat{y}_{ij}$  is the  $i$ th-model prediction for the  $j$ th variable,  $y_{obsj}$  is an observation of the  $j$ th variable,  $\epsilon$  is a small scaling constant,  $\sigma_j$  is the standard deviation of observations of the  $j$ th variable, and  $0 > \delta_A < 1$ , is the constant that describes the threshold probability of adequacy.

This definition states that a model is adequate if there is a high probability that all of its predictions are within the limits of the error distributions for the observations. The scaling factors  $\sigma_j$  allow the level of accuracy required for each prediction variable to be adjusted separately. For example, in a VMA, the value of recommendations for the settings of the ventilator is more sensitive to the partial pressure of oxygen ( $P_aO_2$ ) than it is to the partial pressure of carbon dioxide ( $P_aCO_2$ ). The  $\sigma_j$  for  $P_aO_2$  would be set to a smaller value than

the  $\sigma_j$  for the  $P_aCO_2$ , and models that are relatively less accurate for predicting  $P_aCO_2$  would meet the definition of adequacy.

### Probability of Model Adequacy

As a metric of how well a model fits the observations, I use the posterior probability of model adequacy, given all quantitative observations and qualitative information about the system being modeled. This probability is a measure of how well a model fits the observations. A model that fits the observations well is more likely to make accurate predictions than is a model that fits the observations poorly, and the probability of model adequacy is a measure of the benefit of a model in a control application.

I compute the posterior probability that a model is adequate as

$$\Pr(M_i^A | y) = k_y \Pr(M_i^A) \Pr(y | M_i^A), \quad (3)$$

in which  $k_y = 1 / \Pr(y)$ ,  $\Pr(M_i^A)$  is the prior probability that model  $i$  is adequate, and  $y$  is the vector of quantitative observations of model variables.

Now, consider the event,  $M_i^C$ , that model  $M_i$  is a correct model.\* Because any correct model must also be adequate, it follows that  $\Pr(M_i^A | y) \geq \Pr(M_i^C | y)$ , and

$$\Pr(M_i^A | y) \geq k_y \Pr(M_i^A) \Pr(y | M_i^C). \quad (4)$$

An estimate of a lower bound for the posterior probability of model adequacy is the prior probability of model adequacy times the conditional probability of the observations given that the model is correct.

### Prior Probability of Model Adequacy $\Pr(M_i^A)$

There are many ways that the prior probability of a model being adequate could be assessed. For example, the prior for each model could be assigned by the model builder at the time of model construction. A rule-based expert system could modify the prior distribution according to the context of the model selection. The rule-based approach is analogous to the constraint satisfaction methods previously implemented for model selection [1,9,20].

A more powerful method for assessing  $\Pr(M_i^A)$  is to represent the modeling constraints in a belief network. I implemented a belief network to represent the effects of ten clinical diagnoses on the probability of model

adequacy for the eight physiologic models shown in Figure 1. This belief network takes as inputs the clinical assessment of the physiologic diagnosis and computes a prior probability of model adequacy for the set of alternative models [6,10].

For example, if the node corresponding to the diagnosis of asthma is set to TRUE, the prior probabilities for the models that lack an expanded representation of the airways is decreased.

### Conditional Probability of the Data $\Pr(y | M_i^C)$

Under the assumption that the observation errors are unbiased and normally distributed, the weighted sum of the squared residuals (WSSR) for the fit of a model ( $M_i$ ) to the data ( $y$ ) is distributed as  $\chi^2$  with  $N-p$  degrees of freedom, where  $N$  is the number of data points and  $p$  is the number of fitted parameters [16]. That is,

$$WSSR = \sum_{j=1}^N \left( \frac{|\hat{y}_{ij} - y_{obs,j}|}{\sigma_j} \right)^2 \sim \chi^2_{(N-p)}. \quad (5)$$

$\Pr(y | M_i^A)$  is the likelihood of the model, which is computed directly from the  $\chi^2$  distribution, as  $\Pr(\chi^2_{(N-p)} \geq WSSR)$ . For a correct model, as the number of observations increases, the expected value of the likelihood approaches 0.5, and so the maximum expected value for a model with a prior of 1.0 is 0.5.

In summary, an approximate measure of the benefit of a model is

$$\text{Benefit}(M_i) = \Pr(M_i^A) \text{Min}[0.5, \Pr(y | M_i^A)]. \quad (6)$$

### Local Search in the Graph of Models

I organize a set of models as a GoM, in which the most detailed model is at the top, and successive structural simplifications lead to the least detailed model at the bottom (see Figure 1).

To make a time-dependent model selection, the search algorithm computes dynamically the benefit of each model, and makes a decision after each model fit to evaluate another model in the GoM, or to use the current best model.

The following search algorithm makes dynamic model selections under a time constraint:

1. At  $t = 0$ , consider the least complex model with a prior probability that exceeds a threshold  $(1 - \delta_A)$ .
2. Fit the model to the quantitative observations, which advances the time to  $t = t + t_{model}$ . Compute the ratio of benefit to incremental cost of the additional time needed to make a control recommendation,  $\text{Benefit}(M_i) / (\text{Cost}(t + t_{search}) - \text{Cost}(t))$ .

\*. By "correct" model, I mean a model that produces behavior indistinguishable from that of the system being modeled. As Box said, "All models are wrong—therefore we cannot proclaim a correct one." [3]

3. Examine the adjacent models in the GoM for models that might provide an increase in the net benefit. If the ratio of maximum possible increase in benefit to expected incremental cost exceeds 1, then go to step 2. Otherwise, select the current model, or, if the benefit of the best model available within the time constraint does not meet a minimum criterion of model adequacy, declare no model selection.

After a recommendation is made, and until new observations occur, the search continues for a model with a better fit to the existing data.

### THE GRAPH OF MODELS FOR A VMA

I implemented the search algorithm to perform dynamic selection of models under a time constraint within a GoM of cardiopulmonary models that are suitable for use in a VMA. Each model in the GoM is a series of linked, first order difference equations that describe the circulation of oxygen in the body (see Figure 1). The GoM includes models that vary in their level of detail from a simple, 3-compartment model (the VentPlan model,  $M_{CAV}$ ) to a multicompartment model with detailed airway, circulation and mechanical ventilator components (the VentSim model,  $M_{CAV}$ ).

#### VentPlan: A Simplified Physiologic Model

The VentPlan model is an implementation of the simplified 3-compartment model of oxygen transport [13,18]. For ICU patients with a restricted range of cardiopulmonary abnormalities, however, the VentPlan model has reasonable prediction accuracy [15].

The VentPlan model makes 3 major assumptions that, if violated, may cause the model to make inaccurate predictions:

1. *Airway assumption:* The airway resistance and thoracic compliance are near normal, and have a distribution that is symmetric,
2. *Circulation-compartments assumption:* Any ventilation to perfusion distribution ( $V_A/Q$ ) abnormality is due entirely to a fixed shunting of blood around the lungs, and
3. *Ventilator assumption:* The ventilator's driving pressure can be approximated by a sine wave.

The VentPlan model is unable to represent certain common physiologic abnormalities; the model's assumptions describe the unmodeled abnormalities.

For example, the circulation-compartments assumption describes the fact that perfusion and ventilation occur in a single compartment, and there is a single parameter (the shunt fraction) to explain all blood flow that bypasses areas of gas exchange in the lung. The model makes inaccurate predictions for the oxygen

concentrations in patients with a pulmonary embolus, because these patients have perfused areas of the lung that receive disproportionately less ventilation. A 3-compartment model cannot represent the mismatch of the *distribution* of ventilation and perfusion, and cannot predict the consequences of such mismatch.

Similarly, the lack of a detailed model of the airways (the airway assumption) means that this model is poorly able to make predictions for patients who have abnormal airways. For example, the model makes poor predictions of airway pressures and oxygen concentrations for patients with pulmonary fibrosis, who have markedly decreased thoracic compliance.

#### VentSim: An Expanded Physiologic Model

I initially implemented VentSim as an interactive simulator, to allow a user to explore the effects of alternative ventilator settings for a wide range of simulated patients with varying physiologic abnormalities. VentSim is a ventilator and patient simulation model that expands the VentPlan model and retracts all three VentPlan assumptions. VentSim includes multiple  $V_A/Q$  compartments, an explicit model of distribution of airway resistance and lung compliance, and a detailed model of the mechanical ventilator.

Although VentSim is too computation-intensive to be used at the inner loop of the control algorithm for a VMA, it is a reference model from which I derive intermediate models by applying subsets of the VentPlan simplifying assumptions.

#### Intermediate models in the GoM

The intermediate models in the GoM make either one or two of the VentPlan assumptions, and so they are intermediate in their level of detail and in their level of computation complexity.

For example, model  $M_{CAV}$  differs from  $M_{CAV}$  only in that  $M_{CAV}$  makes the ventilator assumption. This assumption is valid when the inspiration to expiration ratio is close to 0.5, and when there is little asymmetry in the resistance of the airways and compliance of the lungs.

#### Computation Complexity in the GoM

The simplest models in this GoM are solved rapidly using numerical techniques to search for the roots their steady-state equations. The models that contain the expanded ventilator component have no analytic steady-state solution, and must be solved using the more computation-intensive technique of numeric integration.

## Applying Physiologic Models to Quantitative Data

To make a model patient-specific, the model parameters are fitted to the data. When there are few observations, the parameters may be underdetermined.

Empirical Bayes' estimation avoids this problem by adjusting a prior distribution for each model parameter to obtain an approximate posterior parameter distribution in light of the observations. In other words, the fitted parameter values are moved from the modes of their prior distributions only as much as is needed to make the model predict the observations. In the presence of few quantitative observations, the model makes predictions that are typical of patients who are described by the prior distributions [17].

I implemented this method as a modified Levenberg-Marquardt fitting procedure, in which prior distributions on the fitted parameters are treated as if they were observations of the parameters [12,21]. The VentPlan prototype VMA implemented a similar technique [18].

## RESULTS

I generated simulated patient data by setting the parameters of the base model to values that corresponded to varying abnormalities in the airways,  $V_A/Q$  distribution, shunt fraction, and cardiac output. I then added zero-mean, normally distributed noise to these data, and computed the fit of each model to the data.

The results of fitting models  $M_{cav}$  and  $M_{CAV}$  to data from a simulated patient with no significant abnormality, and to data from a simulated patient with a moderately severe airway abnormality and a severe  $V_A/Q$  mismatch are shown in Figure 2. These graphs demonstrate that, for some physiologic states, the least complex model in the GoM makes accurate predictions (Figure 2a), while for other abnormalities, only more complex models are able to explain the observations and make accurate predictions (Figure 2b,2c).

In Figure 3, the benefits are plotted against the costs for the fit of all models to the data shown in Figure 2a. This graph displays the information required to assess the optimal model to select at  $t = 0$ . The plot shows there is wide variation in benefit/cost ratios for the models in this GoM, and shows that  $M_{cav}$  has a benefit/cost advantage over the other models in the GoM, even though  $M_{cav}$  does not provide the highest benefit.

For the data in Figure 2a, the model-selection algorithm began by fitting the simplest model,  $M_{cav}$ . The algorithm then fit  $M_{CAV}$ , and  $M_{CAV}$ , but with no improvement in the benefit (these fits required a total of 5.02 minutes to complete). At that point, the algorithm selected  $M_{cav}$ , since no other model could provide a benefit equal to the expected cost of further computation, even if it provided a perfect fit to the data.

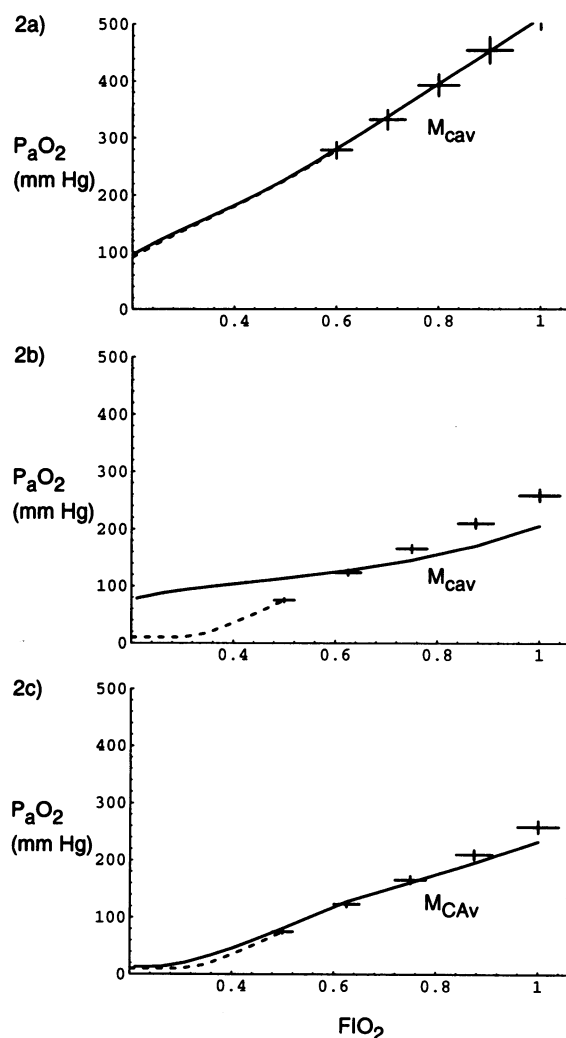


Figure 2. Model fitting to noisy observations of 2 simulated patients. a) Fit of model  $M_{cav}$  (the VentPlan model) to data generated by a simulated patient with no significant abnormality. b) Fit of model  $M_{cav}$  to data generated by a simulated patient with a ventilation-perfusion abnormality and with abnormal airways. c) Fit of model  $M_{CAV}$  to the same data as in b), and using the same set of parameter priors for the fit. Note that the more complex model  $M_{CAV}$  fits the more complex patient more accurately than does  $M_{cav}$ . Legend: solid lines: plot of model with the fitted parameter values; short dashed lines: true behavior of the simulated patient beyond the region of data observations; crosses: simulated patient observations. The vertical length of the crosses is the standard deviation of the normally distributed noise.

For the data in Figure 2b, the algorithm first evaluated  $M_{cav}$ , then  $M_{CAV}$ , and finally  $M_{CAV}$ . These 3 models all fit the data poorly, with benefits of less than  $10^{-44}$ , and the algorithm stopped, declaring no model selection

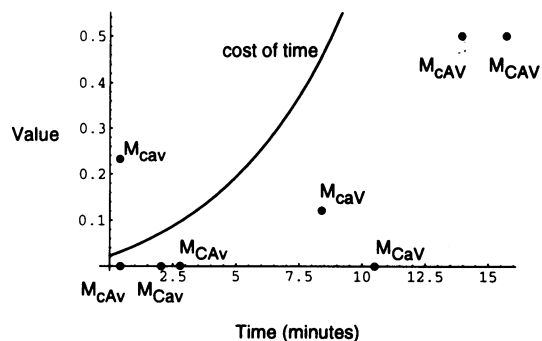


Figure 3. Model benefit versus model cost. For each model fitted to the data shown in Figure 2a), the value of the calculated benefit (Equation 6) is plotted against the cost of the time to fit the model to the data ( $t_{model}$ ). The solid line shows the cost of time function (Equation 1) with  $C=0.1$ ,  $k=5$  minutes and  $t_{search}=1$  minute.

within the time constraint. Interestingly, the next model that would be tested if the time constraint were relaxed, ( $M_{CaV}$ ) represented the best benefit/cost tradeoff, with a benefit of 0.004 and a computation time ( $t_{model}$ ) of 6.5 minutes.

This second case demonstrates the potential impact that the belief network may have, if the network computes prior probabilities  $< 1-\delta_A$  for the less complex models when other information—such as a clinical diagnosis—suggests that less complex models will be inaccurate.

## DISCUSSION

Finding the appropriate assumptions and simplifications that will lead to tractable, yet accurate, models is a difficult task, requiring knowledge of statistical and numerical methods, experience in model building, and expertise in the domain of the application. Traditionally, model-building experts have hand-crafted simulation models of complex domains to meet the required accuracy of a specified task with a minimum of complexity.

Prior implementations of the GoM have been restricted to qualitative or semi-quantitative models in physical domains. This paper presents a methodology for reasoning with sets of continuous models of complex, highly nonlinear systems.

The GoM is a powerful structure to reason with alternative models, but requires that the models be enumerated in advance. An approach to model selection that does not require the prior enumeration of the alternative models involves *composing* models from individually selected components, or submodels [4,9].

Because the rules constraining the selection of model components are a form of simplifying assumptions, compositional modeling is analogous to search in the

GoM. Compositional modeling is a search in a conceptual GoM that consists of all valid combinations of model components.

The compositional approach requires that the components of the model do not have strong interactions—the models must be decomposable. Unfortunately, in the domain of cardiopulmonary physiologic models, the model components interact strongly, and the compositional technique is not directly applicable.

I represent the constraints on model adequacy in a belief network that combines all constraints to compute a prior distribution over the set of available models. This prior distribution focuses the initial search on candidates with no violated constraints. The search metric requires that the system fit individual models to the data to compute a posterior probability of model adequacy. Although the fitting procedure is computationally intensive (especially for the more complex models), the search algorithm ensures that the least complex models that may be adequate are investigated first.

The physiologic models included in the GoM for a VMA are based on relatively simple, and agreed-upon concepts of physiology—the most complex model in the graph does not approach the level of detail of our understanding of the relevant physiologic processes. Models do not have to be particularly sophisticated to be of benefit, however—even linear regression models provide some predictive power in this domain. The simple model in VentPlan provided surprisingly good predictions for a subset of patients in the ICU [14], and the additional levels of detail in any of the more complex models in the VMA GoM should improve the performance of a VMA that incorporates them.

The architecture for the GoM that I propose has, at the top, a *base model*—the most complete and complex model that is implemented [22]. Although this model may not be tractable for a real time application, it is a reference against which the performance of less detailed models can be compared. In my formulation, the less complex models are approximations of the base model that are derived by applying simplifying assumptions to it.

The purpose of reasoning with the GoM is to find a model that is less complex than the base model, but that computes, within the time constraints, predictions that are nearly as good as those of the base model.

The performance evaluation for my dynamic selection of models method asks “If I have a base model, can I select a less complex model that gives a better tradeoff of prediction accuracy and computation time complexity?” How well this model-selection method works is not dependent on whether or not the base model itself

is valid. For this reason, I have not performed an assessment of the validity of the base model in the first phase of this research.

The next phase of this research will include a validation of the prediction accuracy for the selected models, using patient data collected from an online ICU computer charting system, and an assessment of the impact of the belief network on model-search performance.

#### Acknowledgments

I thank Ross Shachter, Lewis Sheiner and Lawrence Fagan for helpful discussions. George Thomsen developed the first VentPlan model and numerous model-solving techniques. Adam Galper graciously provided me with the routines for belief-network evaluation. I especially thank Edward Shortliffe for providing the environment for research in Medical Informatics at Stanford. Preliminary results of this research were presented at the Fourth International Workshop in AI and Statistics, Ft. Lauderdale, FL, on January 6, 1993. This research was supported in part by Grant IRI-9108359 from the National Science Foundation and Grants LM-07033 and LM-04136 from the National Library of Medicine. Computing facilities were provided by the CAMIS Resource, LM-05305.

#### Reference

- [1]. Addanki, S., Cremonini, R. and Penberthy, J.S. Graphs of models. *Artificial Intelligence*, 51:145–177, 1991.
- [2]. Baisingthwaite, J. B. Using computer models to understand complex systems. *The Physiologist* 28:439–442, 1985.
- [3]. Box, G.E.P. *The Collected Works of George E. P. Box*. Belmont, CA: Wadsworth Advanced Books and Software, 1985.
- [4]. Falkenhainer, B. and Forbus, K.D. Compositional modeling of physical systems. *Artificial Intelligence*, 51:95–143, 1991.
- [5]. Horvitz, E. and Rutledge, G. Time-dependent utility and action under uncertainty. *Proceedings of the Seventh Conference on Uncertainty in Artificial Intelligence*, Los Angeles, CA, pp. 151–158. Palo Alto, CA: Morgan Kaufman, July 1991.
- [6]. Lauritzen, S.L. and Spiegelhalter, D.J. Local computations with probabilities on graphical structures and their application to expert systems. *Journal of the Royal Statistical Society*, 50:157–224, 1988.
- [7]. Murthy, S. and Addanki, S. PROMPT: An innovative design tool. *Proceedings of the Sixth National Conference on Artificial Intelligence (AAAI-87)*, Seattle, WA, pp. 637–642, Palo Alto, CA: Morgan Kaufman, August, 1987.
- [8]. Neelankavil, F. *Computer Simulation and Modeling*. Chichester: John Wiley and Sons, 1987.
- [9]. Nayak, P., Addanki, S. and Joskowicz, L. Automated model selection using context-dependent behaviors. *Proceedings of the Tenth National Conference on Artificial Intelligence (AAAI-92)*, San Jose, CA, pp. 710–716, Menlo Park, CA: AAAI Press, July, 1992.
- [10]. Pearl, J. Fusion, propagation and structuring in belief networks. *Artificial Intelligence*, 29:241–248, 1986.
- [11]. Penberthy, S. *Incremental Analysis in the Graph of Models*. S.M. thesis, Department of Computer Science, Massachusetts Institute of Technology, Cambridge, MA, 1987.
- [12]. Press, W.H., Flannery, B.P., Teukolsky, S.A. and Vetterling, W.T. *Numerical Recipes in C: The Art of Scientific Programming*. Cambridge: Cambridge University Press, 1989.
- [13]. Riley, R.L. and Cournand, A. 'Ideal' alveolar air and the analysis of ventilation-perfusion relationships in the lungs. *J. Appl. Physiol.*, 1:825–847, 1949.
- [14]. Rutledge, G.W. Dynamic selection of models under time constraints. *Proceedings of the Second Annual Conference on AI, Simulation and Planning in High Autonomy Systems*, Cocoa Beach, FL, pp. 60–67, Los Alamitos, CA: IEEE Press, April, 1991.
- [15]. Rutledge, G., Thomsen, G., Farr, B., et al. The design and implementation of a ventilator-management advisor. *Artificial Intelligence in Medicine Journal*, 5:67–82, 1993.
- [16]. Schwarz, G., Estimating the dimension of a model. *Ann. Statist.*, 6:461–464, 1978.
- [17]. Sheiner, L.B. and Beal, S.L. Bayesian individualization of pharmacokinetics: simple implementation and comparison with non-Bayesian methods. *Journal of Pharmaceutical Sciences*, 71:1344–1348, 1982.
- [18]. Thomsen, G. and Sheiner, L. SIMV: An application of mathematical modeling in ventilator management. *Proceedings of the Thirteenth Annual Symposium on Computer Applications in Medical Care*, Washington, D.C., pp. 320–324, Los Alamitos, CA: IEEE press, November, 1989.
- [19]. Weld, D. Automated model switching: Discrepancy-driven selection of approximation reformulations. Technical Report 89-08-01, Department of Computer Science, University of Washington, Seattle, WA, October, 1989.
- [20]. Weld, D.S. Reasoning About Model Accuracy. *Artificial Intelligence*, 56:255–300, 1991.
- [21]. Wolfram, S. *Mathematica, A System for Doing Mathematics by Computer*, 2nd ed., Redwood City, CA: Addison Wesley, 1991.
- [22]. Zeigler, B.P. *Theory of Modelling and Simulation*. New York: John Wiley and Sons, 1976.